arXiv:1609.01616v1 [cs.SI] 6 Sep 2016

Private Link Exchange over Social Graphs

Hiep H. Nguyen, Abdessamad Imine, and Michaël Rusinowitch
LORIA/INRIA Nancy-Grand Est, France
Email: {huu-hiep.nguyen,michael.rusinowitch}@inria.fr, abdessamad.imine@loria.fr

***Abstract*—Currently, most of the online social networks (OSN) keep their data secret and in centralized manner. Researchers are allowed to crawl the underlying social graphs (and data) but with limited rates, leading to only partial views of the true social graphs. To overcome this constraint, we may start from user perspective, the contributors of the OSNs. More precisely, if users cautiously collaborate with one another, they can use the very infrastructure of the OSNs to exchange noisy friend lists with their neighbors in several rounds. In the end, they can build local subgraphs, also called local views of the true social graph. In this paper, we propose such protocols for the problem of *private link exchange* over social graphs.**

The problem is unique in the sense that the disseminated data over the links are the links themselves. However, there exist fundamental questions about the feasibility of this model. The first question is how to define simple and effective privacy concepts for the link exchange processes. The second question comes from the high volume of link lists in exchange which may increase exponentially round after round. While storage and computation complexity may be affordable for desktop PCs, communication costs are non-trivial. We address both questions by a simple (α, β)-exchange using Bloom filters.

I. Introduction

Online social networks (OSN) have grown significantly over the last ten years with billions of active users using a variety of social network services. OSNs have revolutionized the way people interact. People join social networking sites to connect and communicate with their friends in real-time. They share interests and activities across political, economic, and geographic borders. As social network sites continue to develop both in number and size, the service providers accumulate unprecedented amount of information about OSN users. As a result, social networks are a valuable data source for research on information societies. In particular, underlying social graphs play a key role in understanding how people form communities, how the OSNs suggest friendship to two users who do not know each other but have many common friends, etc. However, social graphs are not published in clear form due to serious privacy concerns. Instead, they are anonymized in various forms and published to third party consumers such as sociologists, epidemiologists, advertisers and criminologists. Alternatively, social networking sites provide APIs [1] for data crawlers at limited rates and within privacy constraints (e.g. only public friend lists are available). Using this method, the data crawlers can collect friendship information and build a partial (local) view of the target social graph.

To overcome the constraints set by the service providers, we can start from user perspective, i.e. the contributors of OSNs. More precisely, if users cautiously collaborate with one another, they can exchange *noisy* friend lists (containing fake friendships) with their neighbors in several rounds to get better views of the true social graph. Our ideas are based on the fact that user IDs are public (e.g. Facebook profiles are searchable [1]) but the friendships are not so, except when a user leaves his friend list in public mode. Using public user IDs, any user can claim fake links from himself to the users not in his friend list.

The aggregation problem in this paper is unique in the sense that the disseminated data over the links are the links themselves. However, there exist fundamental questions about the feasibility of this model. The first question is how to define simple and effective privacy concepts for the link exchange processes. The second question comes from the high volume of link lists in exchange which may increase exponentially round after round. While storage and computation complexity may not be big problems, communication costs are non-trivial. We address both questions by a simple (α, β)-exchange protocol with or without Bloom filters. To protect true links from inference attacks, we add fake links which are $beta$-fraction of true links. Furthermore, we realize the attenuated propagation of links via the parameter $\alpha \leq 1$.

Basically, we assume that users are *honest-but-curious* (HbC), i.e. they follow the protocol but try to figure out true friendships among noisy friend lists. To preserve link privacy, each node obfuscates its friend list by adding fake links originating from itself to a number of nodes not in its friend list. Then in exchange stage, nodes share only with their friends a fraction of noisy links they possess.

Our contributions are summarized as follows:

- We introduce a novel private link exchange problem as an alternative to social graph crawling and centralized anonymization of data. The problem is distributed and provides a privacy/utility trade-off for all nodes.
- We present two schemes for (α, β)-exchange protocol: Baseline and Bloom filter based. We protect the true links by adding fake links and requiring the propagation probability of links to be attenuated by distance. We analyze the advantages and disadvantages of each scheme.
- We evaluate our proposed schemes on various synthetic graph models and draw a number of critical findings.

The paper is organized as follows. We review the related work for information dissemination in social graphs, distributed anonymization, social trust models and Bloom filter in Section II. Section III briefly introduces key concepts of Bloom filter and our link exchange model. In Section

[1] https://developers.facebook.com/docs/graph-api

IV, we present Baseline (α, β)-exchange that realizes the exchange model by sending noisy link lists in clear form. Section V describes Bloom filter version of (α, β)-exchange with constant complexities and better privacy. We validate the proposed schemes in Section VI. Finally, we present our remarks and suggest future work in Section VII.

Table I summarizes notations used in this paper.

TABLE I: List of notations

Symbol	Definition
$G = (V, E)$	social graph with $N = \|V\|$ and $M = \|E_G\|$
$A(G)$	adjacency matrix of G
D	degree sequence of G (column vector)
$Diam(G)$	diameter of G
$N(u)$	neighbors of node u in G, $d_u = \|N(u)\|$
T	number of exchange rounds
(v, w)	true link between v and w
$(v \rightarrow w)$	fake link generated by v
$L_u(t)$	set of links possessed by u at round t
$L_{uv}(t)$	set of links u sends to v at time t
$\propto$	uniformly at random sampling without replacement
α	fraction of links shared between a pair of nodes
β	fraction of fake links generated at $t = 0$
m	number of bits in Bloom filter
k	number of hash functions used in Bloom filter
n	number of elements in Bloom filter
$Bf_u(t)$	Bloom filter possessed by u at round t
$Bf_{uv}(t)$	Bloom filter u sends to v at time t

II. Related Work

Epidemic spreading [15], [12] is the most related to our work. In [15], Pastor-Satorras et al. study the spreading of infections on scale-free (power-law) networks via the susceptible-infected-susceptible (SIS) model [2]. They find the absence of an epidemic threshold ($\lambda_c = 0$) and its associated critical behavior when the number of nodes goes to infinity using mean-field approximation. Moreno et al. [12] provide a detailed analytical and numerical study of susceptible-infected-removed (SIR) on Watts-Strogatz (WS) small-world model and Barabási-Albert (BA) scale-free model. WS graphs with exponentially distributed degrees can be considered as a *homogeneous* model in which each node hash the same number of links. WS graphs have finite epidemic thresholds. On the contrary, BA graphs with power-law distributed degrees are *heterogeneous* and they expose the weaker resistance to epidemics starting on highly connected nodes.

Giakkoupis et al. [9] introduce a distributed algorithm RIPOSTE for disseminating information in a social network that preserves privacy of nodes. Whenever the information reaches a node, the node decides to either forward the information to his neighbors or drop it. RIPOSTE uses two global parameters δ and λ and satisfies differential privacy by applying Randomized Response Technique (RRT) [6]. Our work is also a form of information dissemination over graphs but it spreads a large number of links, not a single item.

Gossip-based protocols [8], [18] aim at providing alternatives to network-level multicast with good scalability and reliability properties. In these protocols, message redundancy for high reliability is ensured by the fact each member forwards each message to a set of other, randomly chosen, group members. Ganesh et al. [8] propose a fully decentralized and self-configuring protocol SCAMP that provides each member with a partial view of group membership. As the number of participating nodes changes, the size of partial views automatically adapts to the value required to support a gossip algorithm reliably. CYCLON [18] is a protocol for construction of reliable overlay networks. It is targeted to overlays that have low diameter, low clustering, highly symmetric node degrees and highly resilient to massive node failures. These properties belong to random graphs. CYCLON employs enhanced shuffling operation in which nodes select neighbors for cache exchange based on their age.

By exchanging noisy link lists, our schemes are related to distributed graph anonymization [5], [17]. However, rather than producing a single global anonymized graph as in [17], link exchange protocols result in multiple local outputs. In addition, link exchange operates at finest-grained level (node-level) whereas previous works consider a small number of data holders who manage disjoint sets of nodes.

The idea of adding fake links to hide true links appears in a number of earlier studies, e.g. [16], [14]. Shokri et al. [16] propose a method for privacy preservation in collaborative filtering recommendation systems. They develop a model where each user stores locally an offline profile on his own side, hidden from the server, and an online profile on the server from which the server generates the recommendations. Each user arbitrarily contacts other users over time, and modifies his own offline profile through aggregating ratings from other users. The more ratings a user aggregates, the higher privacy he is but lower accuracy in recommendations. Nguyen et al. [14] present a centralized graph anonymization scheme based on edge uncertainty semantics. Fake links are added to probabilistically hide true links. They consider distance-2 fake links to keep higher utility.

III. Preliminaries

In this section, we present the exchange model and attack model. Then we review key concepts about Bloom filter.

A. Exchange Model

We consider a distributed exchange model in which each node possesses his friend list and all nodes participate in the exchange protocol. We work on the following assumptions

- **Assumption 1** The space of node IDs is public. A node can generate fake links to any node. All friend lists (true links) are private, i.e. the existence of true link (u, v) is surely known to u and v only.
- **Assumption 2** A node exchanges messages with its neighbors only. Interacting with neighbors is based on an intuition of trusted relationships: we trust our friends more than any stranger.
- **Assumption 3** A synchronous model is guaranteed by *round-tagged* messages. It means a node prepares the

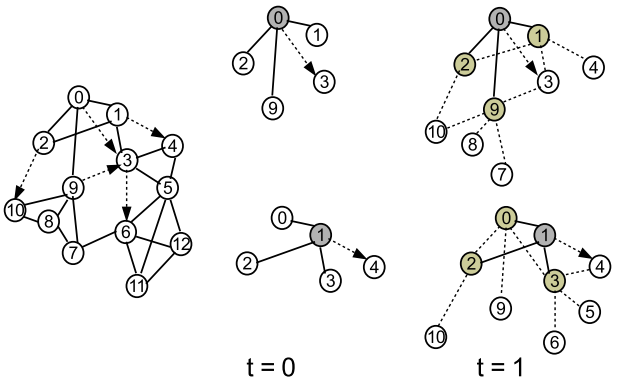

Fig. 1: Link exchange with $\alpha = 1$, $\beta = 1/3$

message for round $t + 1$ if and only if it has received all t-th round messages from his friends.

- **Assumption 4** All nodes are honest-but-curious. They follow the protocol but try to infer true links among noisy links.

Fig. 1 illustrates the exchange model. At round $t = 0$ (initial round), each node u prepare a noisy friend list by adding some fake links (u, v) (i.e. links from u to some people not in his friend list). This is feasible because all user IDs are public (e.g. [1]). For example, node 0 adds a fake link (0,3) and his noisy friend list {(0,1), (0,2), (0,3)} is ready to be exchanged. Similarly, the other nodes prepare their noisy friend lists as in Fig 1. At round $t = 1$, all nodes send and receive noisy friend lists from their neighbors. The local views of nodes 0 and 1 at $t = 1$ are shown in Fig. 1 where the solid lines (resp. the dashed arrows) are the true links (resp. fake links) known by the node and the dashed lines represent noisy links received at the node.

B. Attack Model

We consider honest-but-curious users (nodes) who follow the protocol but try to infer true links among noisy links. We propose a simple inference attack based on frequencies of links arriving to a node. Given a link (v, w) (a true link or a fake link) arriving to node u, if (v, w) does not exist in u's local view, it will be added. Otherwise, its frequency is increased by 1. At the end of the protocol, each node sorts all links in its local view by frequency and selects top K links as true links. How to select the value of K will be discussed later.

By splitting noisy links into two sets of links as above, the inference capability of each node is evaluated by common measures [7]: True Positives (TP), True Negatives (TN), False Positives (FP), False Negatives (FN). As we will see in Section IV, the parameter α introduces an *attenuation effect* for link propagation when $\alpha < 1$. Given a link e, nodes farther from e have lower chance of getting this link. This effect adds another dimension to our privacy model.

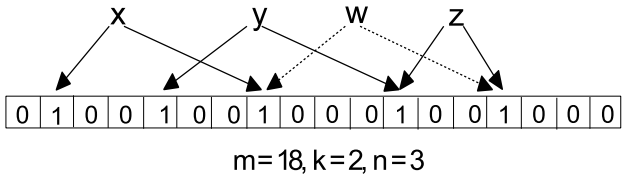

Fig. 2: Bloom filter

C. Bloom Filter

The Bloom filter is a space-efficient probabilistic data structure that supports set membership queries. It was first conceived by Burton Howard Bloom in 1970 [3]. It is used to test whether an element is a member of a set and can result in false positives (claiming an element to belong to the set when it was not inserted), but never in false negatives (reporting an inserted element not in the set).

An empty Bloom filter is an array of m bits, all set to zero. There must also be k different hash functions defined, each of which maps or hashes some set element x to one of the m array positions with a *uniform* random distribution. The number of elements inserted into the Bloom filter is n. Fig. 2 gives an example of Bloom filter with $m = 18$, $k = 2$ and $n = 3$. The MD5 hash algorithm is a popular choice for the hash functions. When an element not in the set w is looked up, it will be hashed by the k hash functions into bit positions. If one of the positions is zero, we conclude that w is not in the set. It may happen that all the bit positions of an element have been set. When this occurs, the Bloom filter will erroneously report that the element is a member of the set, also known as false positives. Fig. 2 shows w as a false positive.

Given the three parameters m, k and n, the false positive probability is (see [4]).

$$p = \left(1 - (1 - \frac{1}{m})^{kn}\right)^k \approx (1 - e^{-kn/m})^k \qquad (1)$$

The false positive probability decreases as m increases or n decreases. Given m and n, the probability of false positives $(1 - e^{-kn/m})^k$ is minimized at $k = k_{opt} = \frac{m}{n} \ln 2$ (see [4]). In this case, the false positive rate $p = (1/2)^k$ or equivalently

$$k = -\log_2 p \qquad (2)$$

IV. Baseline (α, β)-exchange

In this section, we present the main ideas of our proposed (α, β)-exchange and the improvements using Bloom filters.

A. Overview

As shown in Section III-A, the link exchange protocol is straightforward. At the beginning of the protocol, all the nodes agree on the number of rounds T and the two parameters $\alpha \in [0, 1]$, $\beta \geq 0$. Then, each node u prepares his own noisy friend list $L_u(0)$ by setting $L_u(0) = \{(u, v)|v \in N(u)\}$ and adding $\beta N(u)$ fake links in the form $(u \to w)$ where $w \notin N(u)$. At $t = 1$, each node u makes a noisy list $L_{uv}(1)$ for every neighbor v so that $L_{uv}(1)$ contains $\alpha|L_u(0)|$ links sampled from $L_u(0)$. Similarly, node v prepares a noisy list $L_{vu}(1)$

for u. All the nodes sends and receives noisy link lists. Next, each node aggregates noisy link lists by removing duplicate links (if any) and obtains his local view of graph by $L_u(1)$. The round $t = 1$ finishes.

At $t = 2$, the process repeats: all nodes u makes a noisy list $L_{uv}(2)$ for every neighbor v that contains $\alpha|L_u(1)|$ links sampled from $L_u(1)$. They exchange noisy link lists and after receiving all $L_{vu}(2)$ from his friends, node u updates his local view and gets $L_u(2)$. When $t = T$, the protocol terminates.

B. Baseline Scheme

The idea in the previous section is called Baseline (α, β)-exchange as all noisy link lists are in clear form. Algorithm 1 shows steps for Baseline (α, β)-exchange.

Algorithm 1 Baseline (α, β)-exchange

Input: undirected graph $G = (V, E)$, parameters $\alpha \in [0, 1]$, $\beta \geq 0$, number of rounds T
Output: noisy local views of graph $L_u(T), u \in V$

1: // initialization stage
2: **for** $u \in V$ **do**
3: $Fa(u) = \{(u \to w)|w \notin N(u)\}$ s.t. $|Fa(u)| = \beta|N(u)|$
4: $L_u(0) = \{(u, v)|v \in N(u)\} \cup Fa(u)$
5: // exchange stage
6: **for** $t = 1..T$ **do**
7: **for** $(u, v) \in E$ **do**
8: $u : L_{uv}(t) \propto L_u(t-1)$ s.t. $|L_{uv}(t)| = \alpha|L_u(t-1)|$
9: $v : L_{vu}(t) \propto L_v(t-1)$ s.t. $|L_{vu}(t)| = \alpha|L_v(t-1)|$
10: u sends $L_{uv}(t)$ to v
11: v sends $L_{vu}(t)$ to u
12: **for** $u \in V$ **do**
13: $L_u(t) = L_u(t-1) \cup \bigcup_{v \in N(u)} L_{vu}(t)$
return $L_u(T), u \in V$

Given the graph structure $G = (V, E)$, two parameters $\alpha \in [0, 1]$, $\beta \geq 0$ and the number of rounds T. The protocol takes place in two stages. In initialization stage, each node u prepares his own noisy friend list $L_u(0)$ by adding $\beta N(u)$ fake links in the form (u, w) where $w \notin N(u)$ (Lines 3 and 4). In exchange stage (Lines 6-13), at round t, each node u makes a noisy list $L_{uv}(t)$ for every neighbor v that contains $\alpha|L_u(t)|$ links sampled from $L_u(t)$. The exchange happens on every relationship (true link). Each node takes the union of all noisy links he receives before starting the next round.

1) Faster Simulation in A Single PC: For simulation in a single PC, storing all link lists for all nodes in clear form is a costly solution. Moreover, the union operation on lists is time-consuming. We present here a technique to reduce the memory footprint and processing time using bit sets.

Fig. 3 outlines the idea. We have $M' = (1+2\beta)|E_G|$ distinct links consisting of $|E_G|$ true links and $2\beta|E_G|$ fake links. By indexing M' links from 0 to $M' - 1$, the noisy link list at each node is stored in a bit set of size M'. The union of link lists (Line 13 Algorithm 1) is equivalent to an OR operation between bit sets. To prepare $L_{uv}(t)$ for link exchange in round t, node u must recover link IDs in its bit set.

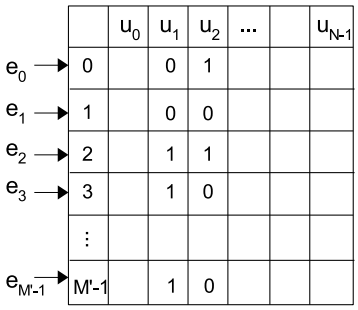

Fig. 3: Fast simulation using bit sets (column vectors)

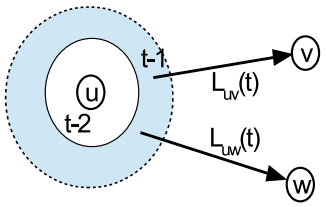

Fig. 4: Incremental volume for $\alpha = 1$

We emphasize that indexing links and storing link IDs in bit sets are only for simulation. In reality, the number of links are unknown to all nodes, so they must run Baseline or Bloom filter (Section V) protocol.

For the case $\alpha = 1$, the exchange volume is reduced further if each node u sends only "new" links, i.e. the links that do not exist in u's list in the previous round. Fig. 4 visualizes this idea in which "new" links are in shaded region and old links are in white region. Note that the incremental volume is valid only for $\alpha = 1$. When $\alpha < 1$, the phenomenon of multipath propagation (Fig. 5) requires both new and old links to be sampled with probability α.

2) Utility-Oriented Initialization: Baseline scheme in the previous section lets a node u generate fake links by connecting u to a certain number of nodes not in u's friend list. This initialization may make local sub graphs at the final round have distorted path distributions due to many fake links connecting faraway nodes. Distorted path distributions reduce the "utility" perceived at each node. Based on the observation of using fake links connecting nearby nodes [14], we suggest a utility-oriented improvement by two-round initialization. We call a fake link $(u \to v)$ *distance-2 link* if $d(u, v) = 2$. For example, $(0 \to 3)$ is a distance-2 fake link while $(2 \to 10)$ is not. Correspondingly, v is called a *distance-2 node* w.r.t u.

We introduce a new parameter $\gamma \in [0, 1]$ which stipulates that each node u create $\gamma\beta d_u$ fake links at $t = 0$ and exchange $\alpha(1 + \gamma\beta)d_u$ randomly chosen links to each of its neighbors. Node u collects node IDs and save them in the set ID_u. At $t = 1$, node u uses node IDs in ID_u to create $(1 - \gamma)\beta d_u$ fake links. Algorithm 2 implements this idea.

The number of distance-2 nodes that u collects in Line 7 of Algorithm 2 is $\alpha(\sum_{v \in N(u)} d_v - d_u - 2Tri(u))$ where $Tri(u)$ is the number of triangles with u as a vertex. Assuming that the set $Fa_0(u)$ contains no distance-2 links (Line 3 Algorithm

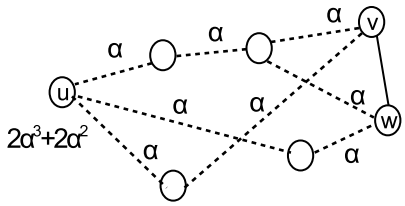

Fig. 5: Multipath link propagation

2). The number of non-distance-2 nodes that u collects is $\sum_{v\in N(u)} \alpha\gamma\beta d_v$. The expected number of distance-2 links that u can create is

$$L2(u) = \frac{(1-\gamma)(\sum_{v\in N(u)} d_v - d_u - 2Tri(u))}{[\sum_{v\in N(u)} d_v - d_u - 2Tri(u)] + \sum_{v\in N(u)} \gamma\beta d_v}$$

$L2(u)$ is a decreasing function of γ. All nodes have the highest (resp. lowest) number of distance-2 fake links at $\gamma = 0$ (resp. $\gamma = 1$). The case of $\gamma = 1$ reduces to standard initialization (Lines 2-4 Algorithm 1).

Algorithm 2 Two-round Initialization

Input: undirected graph $G = (V, E)$, parameters $\alpha, \gamma \in [0, 1]$, $\beta \geq 0$
Output: each node u issues βd_u fake links
1: // t = 0
2: **for** $u \in V$ **do**
3: $Fa_0(u) = \{(u \to w)|w \notin N(u)\}$ s.t. $|Fa_0(u)| = \gamma\beta|N(u)|$
4: $L_u(0) = \{(u, v)|v \in N(u)\} \cup Fa_0(u)$
5: // t = 1
6: **for** $(u, v) \in E$ **do**
7: u and v exchange α-fraction of their links
8: **for** $u \in V$ **do**
9: u aggregates all links it knows into $L_u(1)$
10: $ID_u = \{w|w = v_1 \wedge w = v_2, (v_1, v_2) \in L_u(1)\}\backslash\{u, N(u)\}$
11: $Fa_1(u) = \{(u \to w)|w \in ID_u\}$
12: s.t. $|Fa_1(u)| = (1-\gamma)\beta|N(u)|$
13: $L_u(1) = L_u(1) \cup Fa_1(u)$

C. Complexity Analysis

Let A be the adjacency matrix of G and D be the column vector of degree sequence of nodes, the number of links at all nodes is upper bounded by the following vector, where I_N is the identity matrix of size N.

$$LU(t) = (I_N + \alpha A)^t(1 + \beta)D \tag{3}$$

We say $LU(t)$ is an "upper-bound" because $LU(t)$ accepts duplicate links. More precisely, let $LU_u(t)$ and $LU_{uv}(t)$ be the noisy link lists at node u and for exchange without removing duplicate links as in Line 13 Algorithm 1. We have $LU_u(t) = LU_u(t-1) + \sum_{v\in N(u)} LU_{vu}(t)$, where "+" denotes *multiset* semantics. Clearly, $L_u(t) < LU_u(t)$.

Note that the number of rounds T can be small because of the following analysis. We have four simple facts (see Fig. 5)

1) a true link (v, w) is propagated to node u at round t if and only if $\min\{d(u, v), d(u, w)\} = t$ for $\alpha = 1$.
2) a fake link $(v \to w)$ is propagated to node u at round t if and only if $d(u, v) = t$ for $\alpha = 1$.
3) a true link (v, w) is propagated to node u at round t with probability $\sum_{p_l\in P(u,v)\cup P(u,w)} \alpha^l$ for $\alpha < 1$. Here p_l is a path of length l from u to v or w.
4) a fake link $(v \to w)$ is propagated to node u at round t with probability $\sum_{p_l\in P(u,v)} \alpha^l$ for $\alpha < 1$.

We consider three cases.

Case 1: $\alpha = 1, \beta = 0$ In this case, there is no fake links. Using Fact 1, we have $|L_u(Diam(G) - 1)| = m$, i.e. every node u receives all true links in G after $(Diam(G)-1$ rounds.

Case 2: $\alpha = 1, \beta > 0$ In this case, there are $2\beta m$ fake links. Using Facts 1 and 2, we have $|L_u(Diam(G))| = (1 + 2\beta)m$, i.e. every node u receives all true links and fake links in G after $Diam(G)$ rounds.

Case 3: $\alpha < 1, \beta \geq 0$ In this case, there are $2\beta m$ fake links. Using Facts 3 and 4, every node u receives all true links (v, w) in G after T rounds if

$$\sum_{t=1}^{T} [(\alpha A)^t]_{vu} + [(\alpha A)^t]_{wu} \geq 1 \tag{4}$$

and all fake links $(v \to w)$ if

$$\sum_{t=1}^{T} [(\alpha A)^t]_{vu} \geq 1 \tag{5}$$

The protocol's complexity is measured in storage, computation and communication. Because all links are stored in clear form, all complexities increase round by round (except the trivial case $\alpha = 0$). They are also upper bounded by the total links in graph, which is $(1 + 2\beta)|E_G|$. Intuitively, low-degree nodes will incur lower complexities than high-degree nodes. However, as t increases, the gap gets narrower. In Section V, we will achieve constant complexities by using Bloom filters.

D. Privacy Analysis

In this section, we discuss the link inference attacks that can be mounted by nodes. Each node has knowledge about the true links connecting itself to its neighbors and the fake links it creates before the first round as well as the fake links pointing to it. The remaining links (denoted as B_u) stored at node u are subject to an inference attack by u. As discussed in Section III-B, u may mount an inference attack by sorting links in B_u by weight and picks top-K links as true links.

In Baseline (α, β)-exchange, the ratio of true links over fake links is $\frac{1}{\beta}$. Each user, therefore, can set $K = \frac{|B_u|}{1+\beta}$ and divide B_u into two sets T_u (predicted true links) and F_u (predicted fake links). The numbers of true positives, true negatives, false positives and false negatives are (see Fig. 6 for an illustration)

$$TP_u = |E_G \cap T_u|\ , FP_u = |T_u \setminus E_G| \tag{6}$$

$$FN_u = |E_G \cap F_u|\ , TN_u = |F_u \setminus E_G| \tag{7}$$

The precision, recall and F1 score are defined as $Prec = TP_u/(TP_u + FP_u)$, $Recall = TP_u/(TP_u + FN_u)$ and $F1 = 2 * Prec * Recall/(Prec + Recall)$.

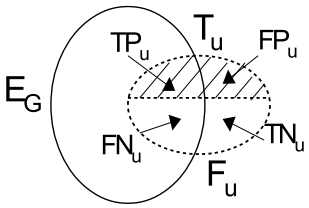

Fig. 6: Inference attack measures

V. Bloom Filter Based Scheme

A. Motivation

Baseline (α, β)-exchange has several drawbacks that motivate a better approach. First, all link lists are in clear form, allowing nodes to store link frequencies for inference attack (Section IV-D). If we obfuscate link lists, this kind of attack may be mitigated. Hashing could be a solution. Second, sending link lists in clear form may incur high communication cost, especially at high degree nodes. Assuming that all node IDs are in range $\{0...2^{32}-1\}$, i.e. each ID needs 4 bytes, each link is encoded in 8 bytes. Given a link list, a better way to encode it is to store all links (u, v_i) incident to u by $\{u|\{v_i\}\}$. In this way, the message length for a link list can be reduced up to 50%. In average, each link costs between 32-bit to 64-bit. Using Bloom filters, the number of bits per link may be reduced. For example, with $k = 4$, the number of bits per link is $k/\ln 2 \approx 5.8$.

This section introduces a Bloom filter based approach. Compared to Baseline approach, it has several advantages and limitations. Bloom filters, by encoding links in compact forms, reduce the storage and communication costs. The computation at each node is also much simpler thanks to logical OR operation compared to set unions in Baseline.

B. Bloom Filter Based Scheme

Algorithm 3 describes steps of Bloom filter version of (α, β)-exchange. As for inputs, we add a global false positive probability p and the number of links $|E_G|$. As analyzed in [4], the number of hash functions k is set to $\lceil -\log_2 p \rceil$ (Line 2). The number of bits per link is $c = k/\ln 2$ (Line 3). The length of every Bloom filter is $m = c.|E_G|$ (Line 4). Then, each node u initializes its Bloom filter $Bf_u(0)$ by hashing all links in $L_u(0)$ using k hash functions. At the same time, all nodes send their noisy links $L_u(0)$ to the coordinator who will gather all links into the list L. This list will be used in the recovery stage.

In the exchange stage, each pair of nodes (u, v) prepare and exchange noisy link lists in encoded form $Bf_{uv}(t)$ and $Bf_{vu}(t)$ (Lines 14-18). Before the next round, each node aggregates all Bloom filters sent to it by taking the OR operation. (Lines 19 and 20). Finally, the recovery stage helps each node to obtain its noisy local view $L_u(T)$. In this stage, the coordinator sends to L to all nodes. If we omit the role of the coordinator (Lines 5,11 and 23), each node u has to try hash $\frac{N(N-1)}{2}$ possible links against its final Bloom filter $Bf_u(T)$.

Algorithm 3 Bloom filter (α, β)-exchange

Input: undirected graph $G = (V, E)$, parameters $\alpha \in [0, 1]$, $\beta \geq 0$, number of rounds T, false positive probability p
Output: noisy local views of graph $L_u(T), u \in V$

1: // initialization stage
2: $k = \lceil -\log_2 p \rceil$ (see equation (2))
3: $c = k/\ln 2$
4: $m = c.|E_G|$
5: $L = \emptyset$
6: **for** $u \in V$ **do**
7: $Bf_u(0) =$ BloomFilter(k,m,c)
8: $Fa(u) = \{(u \to w)|w \notin N(u)\}$ s.t. $|Fa(u)| = \beta|N(u)|$
9: $L_u(0) = \{(u, v)|v \in N(u)\} \cup Fa(u)$
10: Hash all $e \in L_u(0)$ into $Bf_u(0)$
11: $L = L \cup L_u(0)$
12: // exchange stage
13: **for** $t = 1..T$ **do**
14: **for** $(u, v) \in E$ **do**
15: u prepares $Bf_{uv}(t) =$ BitErasure$(Bf_u(t-1), \alpha)$
16: v prepares $Bf_{vu}(t) =$ BitErasure$(Bf_v(t-1), \alpha)$
17: u sends $Bf_{uv}(t)$ to v
18: v sends $Bf_{vu}(t)$ to u
19: **for** $u \in V$ **do**
20: $Bf_u(t) = Bf_u(t-1) \vee \bigvee_{v \in N(u)} Bf_{vu}(t)$
21: // link recovery stage
22: **for** $u \in V$ **do**
23: $L_u(T) =$ Hash$(L, Bf_u(T))$
return $L_u(T), u \in V$

1) Bit Erasure: Because Bloom filters store links information in encoded form, we have to simulate the α-sampling steps (Lines 8 and 9, Algorithm 1).

α-sampling is equivalent to "deletion" of $(1-\alpha)|Bf_u(t-1)|$ elements from $Bf_u(t-1)$. We can perform this operation by recovering elements in $Bf_u(t-1)$ then explicitly keeping α-fraction of elements and hashing these elements to an empty Bloom filter. This approach, however, is costly because the node must try $\frac{N(N-1)}{2}$ possible links. As a result, an implicit removal of $(1-\alpha)$-fraction of elements is needed.

Resetting one bit causes one or several misses (false negatives) and possibly reduces false positives. For example, resetting the second bit in Bloom filter (Fig. 2) makes x a false negative whereas resetting the 12th-bit makes both y and z disappear. Moreover, if the 8-th bit is reset, x becomes a false negative and w is no longer a false positive.

Let m_1 be the number of 1-bits in Bloom filter $Bf_u(t-1)$ and s be the number of randomly reset bits ($s < m_1$), the probability of a true positive remaining in Bloom filter is

$$(1 - \frac{s}{m_1})^k \tag{8}$$

If omitting the effect of false positives (which is reduced as illustrated above), the formula (8) is exactly the sampling fraction α. In other words,

$$\alpha = (1 - \frac{s}{m_1})^k \Rightarrow s = m_1(1 - \alpha^{1/k}) \tag{9}$$

We can see that s is a decreasing function of α and k. An illustration of this fact is shown in Fig. 7.

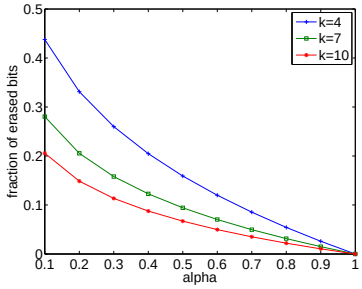

Fig. 7: Fraction of erased bits as a function of α and k

Algorithm 4 realizes α-sampling implicitly via bit erasure.

Algorithm 4 Bit Erasure

Input: Bloom filter B, parameter $\alpha \in [0, 1]$, number of hashes k
Output: Bloom filter B' that contains approximately α fraction of elements in B

1: $B' = B$
2: $M_1 = \{i | B(i) = 1\}$
3: $m_1 = |M_1|$
4: $s = \lfloor m_1(1 - \alpha^{1/k} \rfloor$
5: randomly reset s bits in m_1 positions of B'
6: **return** B'

2) Bloom Filter Compression: In Algorithm 3, all Bloom filters stored at nodes and transmitted between nodes are of length m bits where $m = |E_G|k/\ln 2$. For $p = 0.1$, we have $k = 4$ and $m \approx 5.8|E_G|$. For $p = 0.01$, we have $k = 7$ and $m \approx 10.1|E_G|$. For million-scale graphs with hundreds of millions of links, the length of Bloom filters would be hundreds of megabytes. This is undesirable for message transmission although storage and computation are not big problems. However, we observe that as in Baseline (α, β)-exchange, not all messages have the length of $\Theta(E_G)$. Thus, lossless data compression is a useful tool for Bloom filter exchange.

Arithmetic coding [11] is such a lossless compression scheme. Arithmetic coding differs from other forms of entropy encoding, such as Huffman coding [10]. Huffman coding separates the input into component symbols with symbol probabilities approximated by negative powers of two and replaces each with a code. Arithmetic coding encodes the entire message into a single number, a fraction f where $0.0 \leq f < 1.0$.

C. Complexity and Privacy Analysis

Thanks to constant sizes of bit arrays and constant time for OR operations, the total communication cost of Bloom Filter scheme is constant and the aggregation of noisy link lists is constant too. However, Bloom Filter scheme incurs an extra recovery step at all nodes. Each node needs to download the full noisy link set L from the coordinator. As we confirm in Section VI-B, the exchange time of Bloom Filter scheme is much lower than that of Baseline, but the recovery step costs higher time complexity.

As mentioned in Section V-A, all link lists are obfuscated in Bloom filters, frequency-based inference attacks may be mitigated if the set of all links L is revealed to all nodes only after the final round. The ratio of true links over fake links in Bloom Filter scheme is almost identical to that of Baseline. The reason lies in the independence of all links in exchange protocols. All links have the same probability to be sampled and sent to neighbors of nodes. Interestingly, Bloom Filter helps reduce the true/fake link ratio faster than Baseline for small α (Section VI-A) thanks to its inherent false positives as well as false negatives caused by bit erasure.

VI. Evaluation

In this section, we empirically evaluate the performance of our proposed schemes on synthetic graphs. All algorithms are implemented in Java and run on a desktop PC with $Intel^{®}$ Core i7-4770@ 3.4Ghz, 16GB memory.

Two kinds of synthetic graphs are generated: Barabási-Albert power-law (PL) graphs and Erdös-Rényi (ER) random graphs [13]. Table II lists six synthetic graphs used in our experiments. Each test case is run 10 times. We abbreviate the two schemes Baseline (BS) and BloomFilter-based (BF).

We choose $\alpha \in \{0.25, 0.5, 0.75, 1.0\}$ and $\beta \in \{0.5, 1.0\}$. The default number of hash functions k is 4.

TABLE II: Synthetic graphs

Graph	#Nodes	#Links	Diameter
PL1	10,000	29,990	7
PL2	**10,000**	**49,970**	**6**
PL3	10,000	99,872	5
ER1	10,000	30,076	10
ER2	**10,000**	**50,424**	**7**
ER3	10,000	99,615	5

A. Message Volume and Inference Attacks

We investigate the message volume by the total number of true/fake links at all nodes after each round $t = 1..Diam(G)$. These values are normalized by dividing them by $N.M.(1 + 2\beta)$. We also estimate the inference attacks by the ratio between the number of true links and the number of fake links. Fig. 8 and Fig. 9 show two-y-axis charts. The left y-axis is for the normalized number of links. The right y-axis is for the ratios.

Several observations can be made clearly from Figures 8 and 9. First, the number of true/fake links increases exponentially and converges fast as all nodes reach the round at $Diam(G)$. For $\alpha = 0.25$, Baseline does not converge because not all links are propagated to all nodes. Bloom filter scheme produces higher number of true/fake links, especially at $\alpha = 0.25, 0.5$. For larger values of α, the two schemes almost coincide. Second, the ratio of true links over fake links decreases round by round and converges to $\frac{1}{2\beta}$. In early rounds, the ratios are lower than $\frac{1}{\beta}$. Higher the ratio, higher inference risk of true links. Clearly, Bloom Filter scheme reduces the risk better than Baseline for $\alpha = 0.25, 0.5$ in later rounds.

Fig. 10 displays the distribution of link volume collected at sample nodes. We sort V by degree and take 100 sample

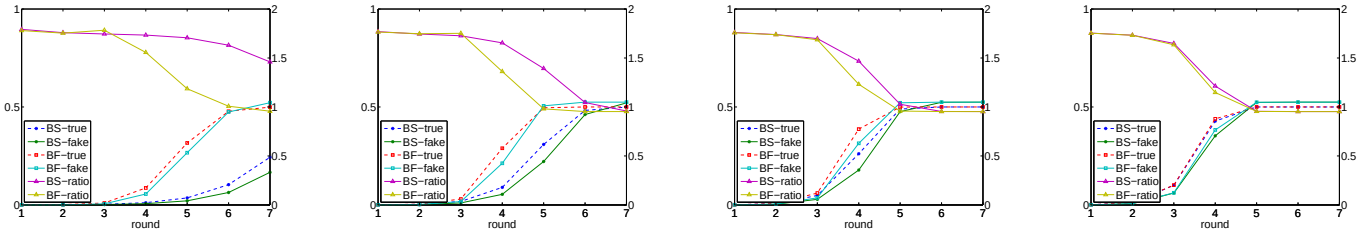

(a) $\alpha = 0.25, \beta = 0.5$ (b) $\alpha = 0.5, \beta = 0.5$ (c) $\alpha = 0.75, \beta = 0.5$ (d) $\alpha = 1.0, \beta = 0.5$

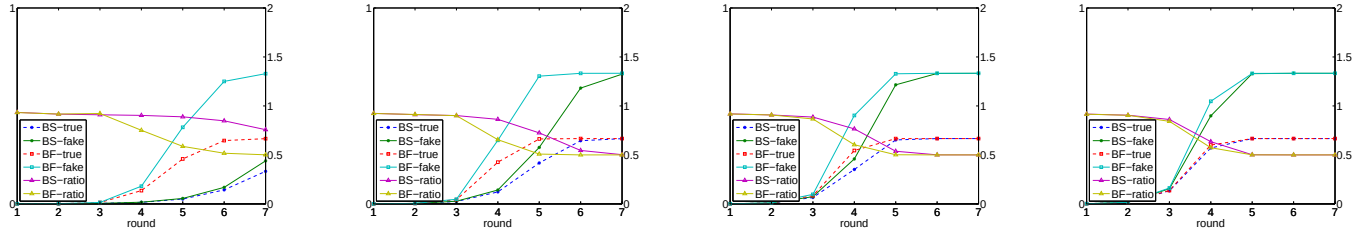

(e) $\alpha = 0.25, \beta = 1.0$ (f) $\alpha = 0.5, \beta = 1.0$ (g) $\alpha = 0.75, \beta = 1.0$ (h) $\alpha = 1.0, \beta = 1.0$

Fig. 8: Normalized number of true/fake links and link ratios on ER2

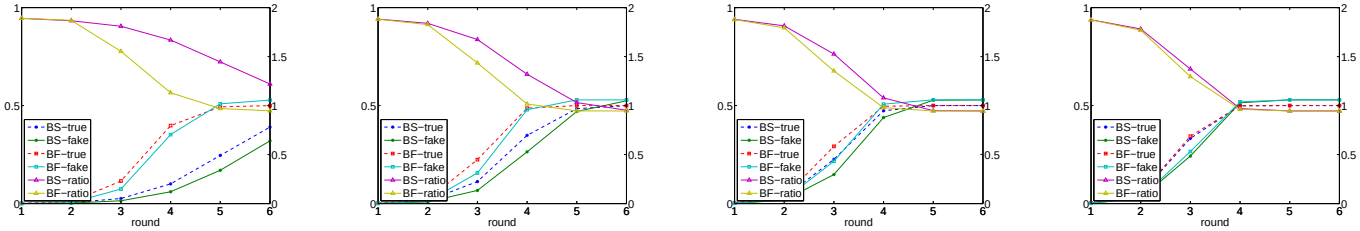

(a) $\alpha = 0.25, \beta = 0.5$ (b) $\alpha = 0.5, \beta = 0.5$ (c) $\alpha = 0.75, \beta = 0.5$ (d) $\alpha = 1.0, \beta = 0.5$

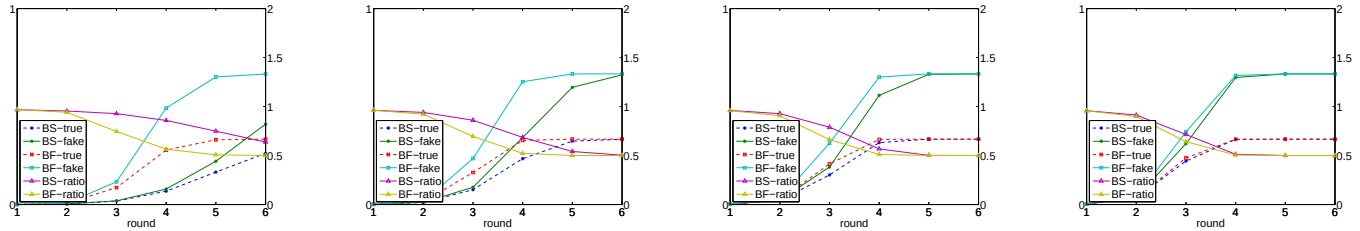

(e) $\alpha = 0.25, \beta = 1.0$ (f) $\alpha = 0.5, \beta = 1.0$ (g) $\alpha = 0.75, \beta = 1.0$ (h) $\alpha = 1.0, \beta = 1.0$

Fig. 9: Normalized number of true/fake links and link ratios on PL2

nodes. ER graphs which are commonly called *homogeneous* graphs show nearly uniform distributions for various values of (α, β). On the contrary, PL graphs are *heterogeneous* ones and sample nodes exhibit much more random distributions.

The inference attack on Baseline scheme (Section IV-D) is shown in Fig. 11. The average F1 scores for two values of β are plotted at different rounds of Baseline protocol. We observe that the scores are quite close to the theoretical values $1/(1+\beta)$ (see the dashed lines). On ER2 graph, the inference attack is more effective at latter rounds and for larger α while this is not clear on PL2.

B. Bloom Filter Scheme

In this section, we examine the performance of Bloom Filter scheme. We set the false positive rate of Bloom Filter to 0.1, 0.01 and 0.001 (the number of hash functions k is 4,7 and 10 respectively). Fig. 12 displays the normalized number of true/fake links by Baseline and Bloom Filter with different false positive rates. We find that lower false positive rates make no improvement for $\alpha = 0.25, 0.5$. Bit Erasure (Algorithm 4) causes this effect. Lower α means more bits to be erased in Bloom filters. Consequently, the number of false positive links and false negative links is amplified round by round for small α.

We compare the communication complexity of Baseline and Bloom Filter schemes. Fig. 13 reports the number of bytes transmitted among nodes after each round in Baseline and Bloom Filter (with or without compression). Baseline scheme stores links in clear form, so it incurs exponential communication complexity. As discussed in Section V-A, we assume that each node ID cost 4 bytes and a link list of length l may be stored compactly in $4l$ bytes. Bloom Filter uses constant-sized bit arrays, so its communication cost is constant too. However, each node running Bloom Filter scheme has to download the full noisy list of $(1 + 2\beta)M$ links after the final round to find which links are contained in its bit array.

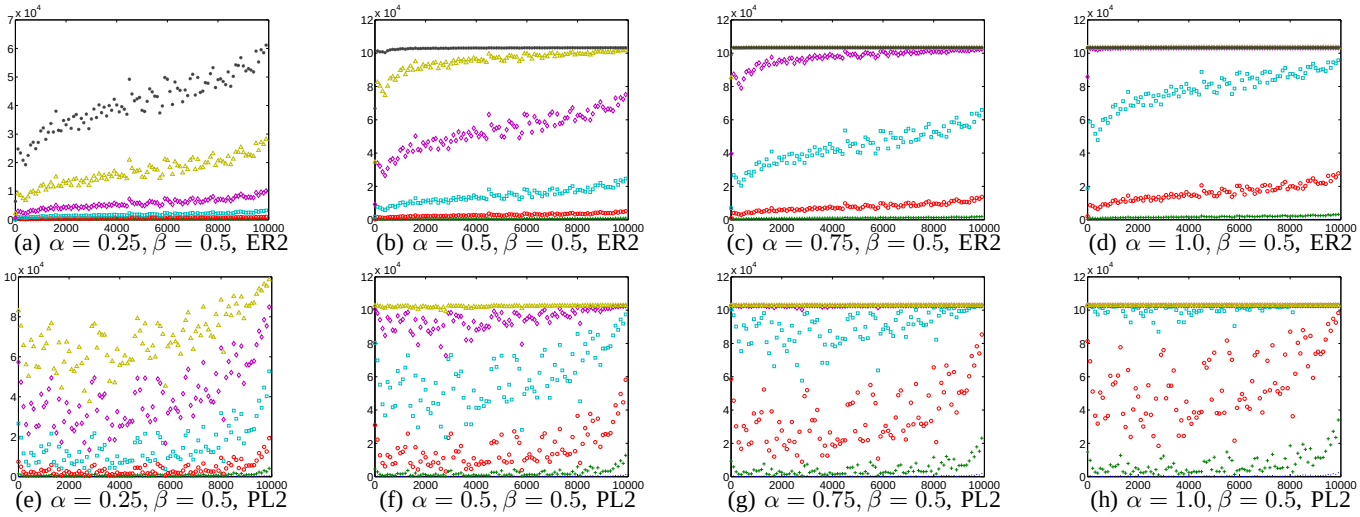

(a) $\alpha = 0.25, \beta = 0.5$, ER2 (b) $\alpha = 0.5, \beta = 0.5$, ER2 (c) $\alpha = 0.75, \beta = 0.5$, ER2 (d) $\alpha = 1.0, \beta = 0.5$, ER2

(e) $\alpha = 0.25, \beta = 0.5$, PL2 (f) $\alpha = 0.5, \beta = 0.5$, PL2 (g) $\alpha = 0.75, \beta = 0.5$, PL2 (h) $\alpha = 1.0, \beta = 0.5$, PL2

Fig. 10: Number of links at sampled nodes ($t = 1(.), t = 2(+), t = 3(\circ), t = 4(\square), t = 5(\diamond), t = 6(\triangle), t = 7(*)$)

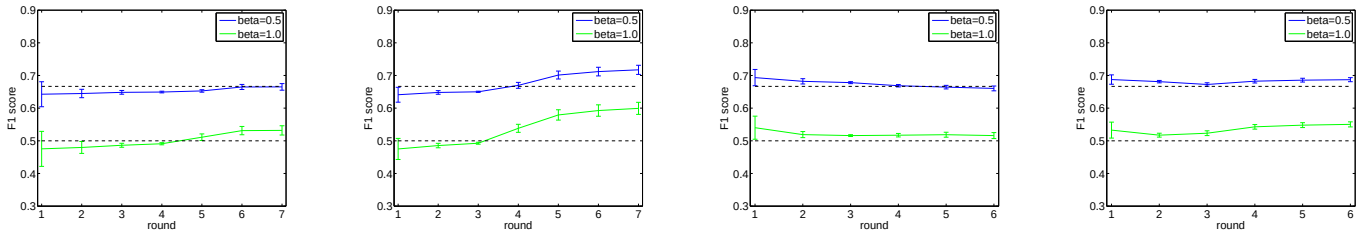

(a) ER2, $\alpha = 0.5$ (b) ER2, $\alpha = 1.0$ (c) PL2, $\alpha = 0.5$ (d) PL2, $\alpha = 1.0$

Fig. 11: Inference attacks

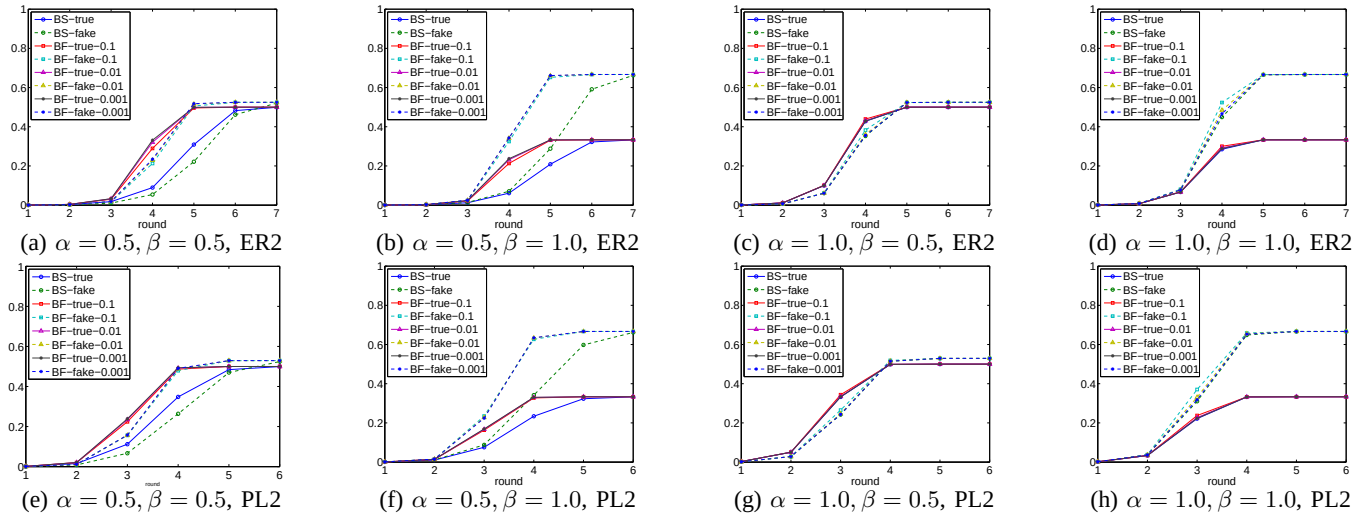

(a) $\alpha = 0.5, \beta = 0.5$, ER2 (b) $\alpha = 0.5, \beta = 1.0$, ER2 (c) $\alpha = 1.0, \beta = 0.5$, ER2 (d) $\alpha = 1.0, \beta = 1.0$, ER2

(e) $\alpha = 0.5, \beta = 0.5$, PL2 (f) $\alpha = 0.5, \beta = 1.0$, PL2 (g) $\alpha = 1.0, \beta = 0.5$, PL2 (h) $\alpha = 1.0, \beta = 1.0$, PL2

Fig. 12: Normalized number of true/fake links by different false positive rates

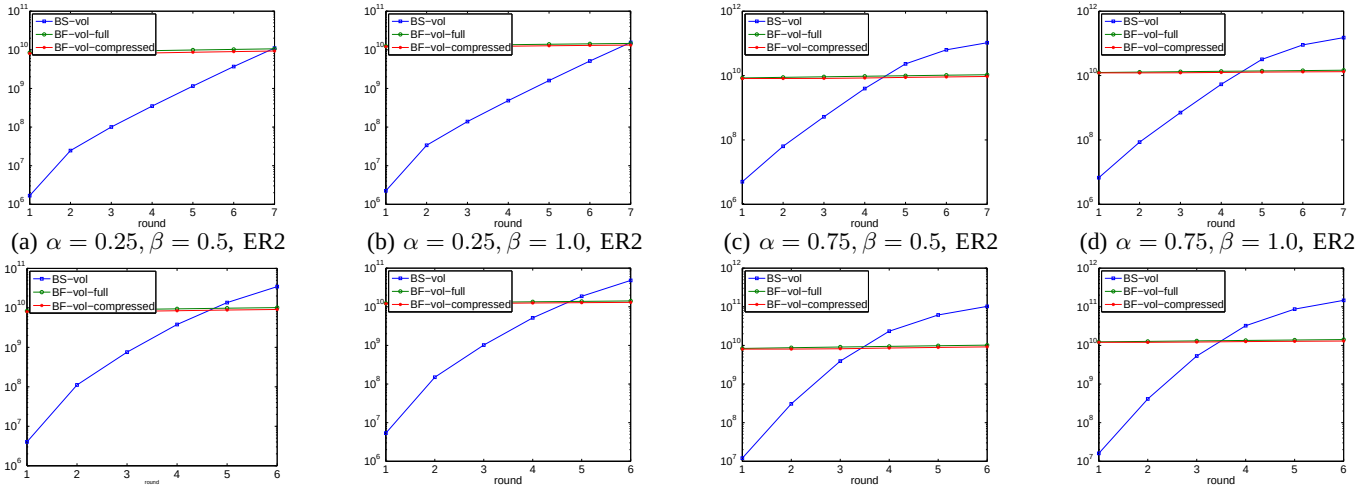

(a) $\alpha = 0.25, \beta = 0.5$, ER2 (b) $\alpha = 0.25, \beta = 1.0$, ER2 (c) $\alpha = 0.75, \beta = 0.5$, ER2 (d) $\alpha = 0.75, \beta = 1.0$, ER2

(e) $\alpha = 0.25, \beta = 0.5$, PL2 (f) $\alpha = 0.25, \beta = 1.0$, PL2 (g) $\alpha = 0.75, \beta = 0.5$, PL2 (h) $\alpha = 0.75, \beta = 1.0$, PL2

Fig. 13: Communication complexity. Y-axis is the total number of bytes transmitted among nodes (log-scale)

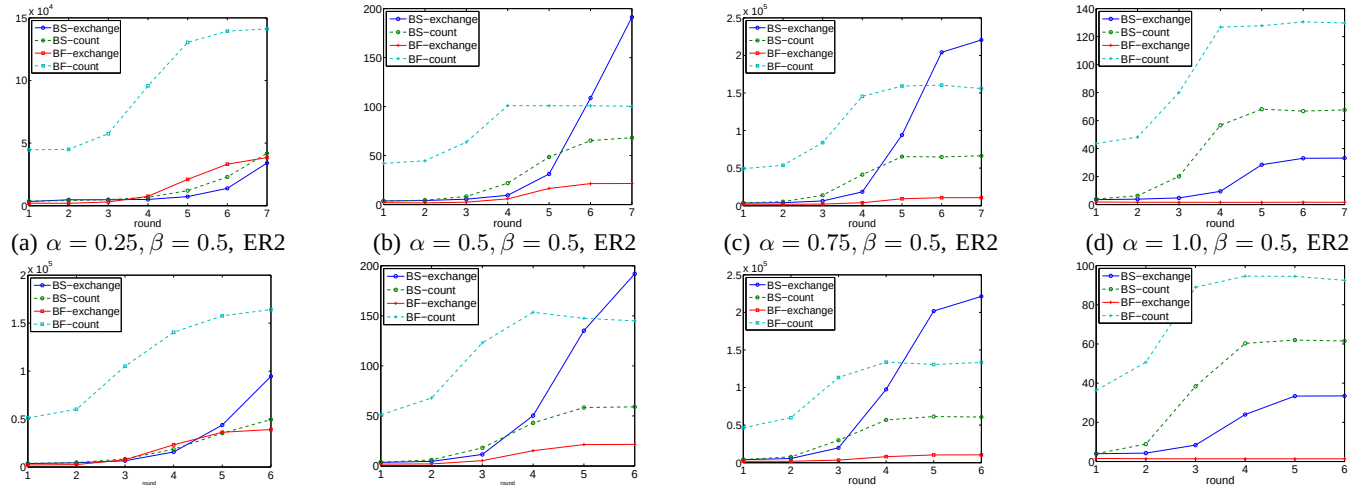

(a) $\alpha = 0.25, \beta = 0.5$, ER2 (b) $\alpha = 0.5, \beta = 0.5$, ER2 (c) $\alpha = 0.75, \beta = 0.5$, ER2 (d) $\alpha = 1.0, \beta = 0.5$, ER2

(e) $\alpha = 0.25, \beta = 0.5$, PL2 (f) $\alpha = 0.5, \beta = 0.5$, PL2 (g) $\alpha = 0.75, \beta = 0.5$, PL2 (h) $\alpha = 1.0, \beta = 0.5$, PL2

Fig. 14: Total simulation runtime of all nodes (in millisecond)

The number of bytes in the download step is $4N(1 + 2\beta)M$ bytes for N nodes. The communication cost of the download step dominates that of bit array exchange. Using bit array compression (Section V-B2), Bloom Filter scheme reduces the message size a little bit, especially at early rounds when a large part of bit arrays are zero bits. For $\alpha = 0.75$, Bloom Filter scheme saves the communication cost in the last three rounds in both ER2 and PL2. For $\alpha = 0.25$, it is worse than Baseline in all rounds (except the final round) on ER2 and in the first four rounds on PL2.

In Fig. 14, we compare the runtime of Baseline and Bloom Filter simulations in a single PC. In each round, each node updates its link set (`count` operation) by aggregating noisy link lists from its neighbors. Then, each node prepares (`exchange` operation), for the next round, new noisy link lists sampled from its link set. At $\alpha < 1$, the exchange operations cost an increasing time as more rounds are considered. Higher α makes the link sampling slower. Only at $\alpha = 1$, we have fast exchange operations. In particular, the exchange runtime of Bloom Filter scheme is constant for $\alpha = 1$ and is an increasing function of round for $\alpha < 1$ due to bit erasure operations. The count operation of Bloom Filter dominates that of Baseline because each node has to hash the full link set to recover its noisy link set at each round.

C. Utility-Oriented Initialization

In this section, we illustrate the benefit of two-round initialization (Algorithm 2). We set $\gamma = 0.0, 0.5$ and denote the enhanced scheme as *D2*. Several utility metrics are chosen as follows.

- Power-law exponent of degree sequence: PL is the estimate of η assuming the degree sequence follows a power-law $n_d \sim d^{-\eta}$ where n_d is the number of d-degree nodes.

- Clustering coefficient: $CC = \frac{3N_\Delta}{N_3}$ where N_Δ is the number of triangles and N_3 is the number of connected triples.

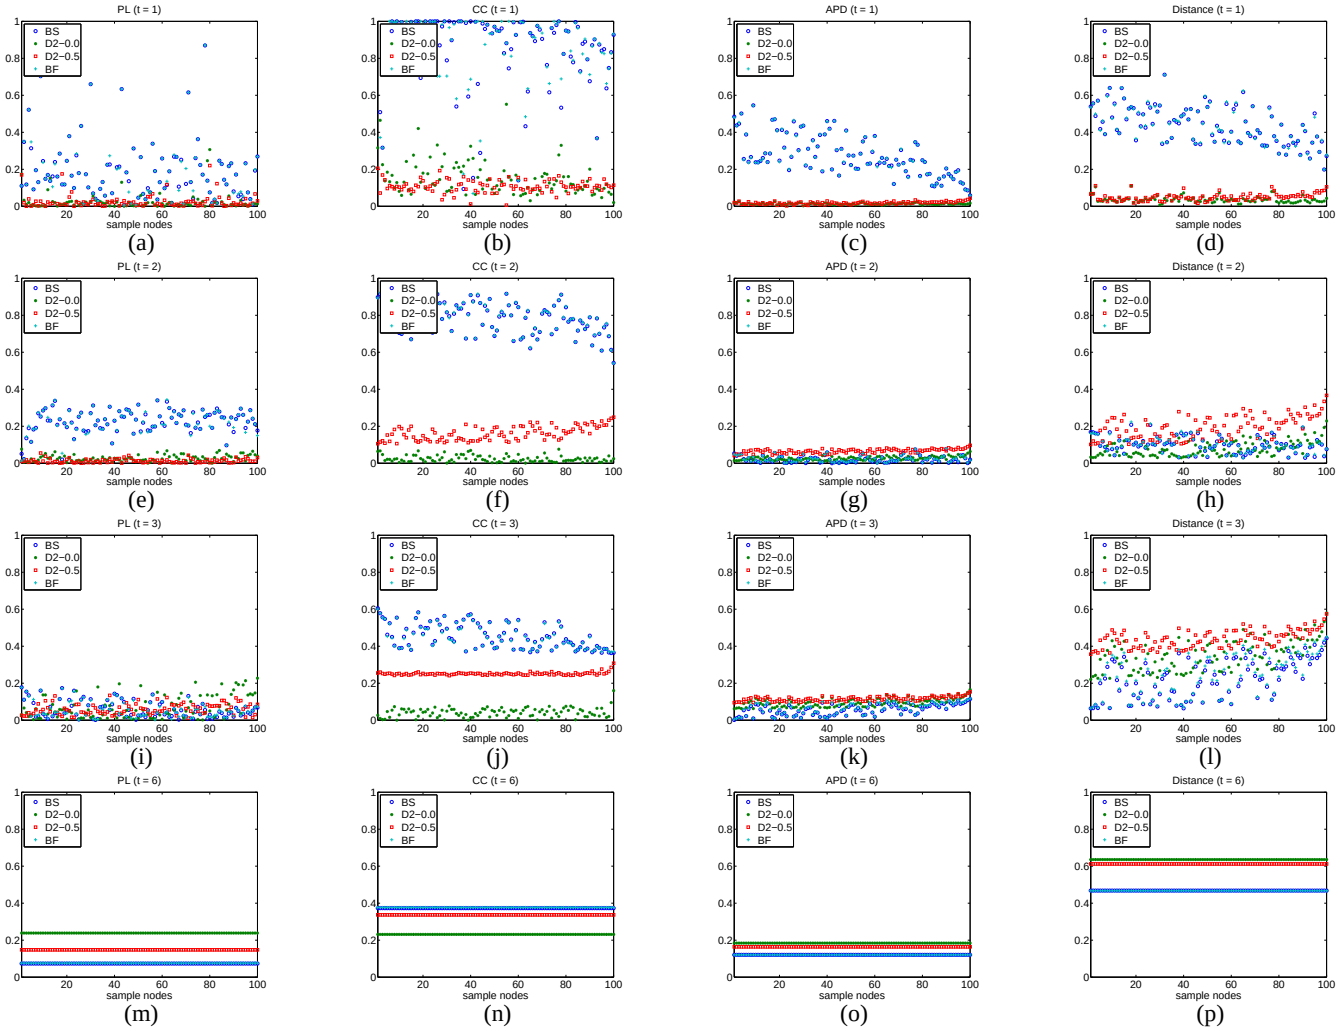

Fig. 15: Utility relative errors on PL2 ($\alpha = 1.0, \beta = 0.5$)

- Average distance: APD is the average distance among all pairs of vertices that are path-connected.

- Distance distribution: $Distance$ is the normalized node-pair shortest-path histogram.

We take 100 sample nodes by degree and compare local aggregated graphs to the ground truth. The ground truth is computed by setting $\beta = 0$ in Baseline scheme. Fig. 15 shows the benefit of two-round initialization (D2-0.0 and D2-0.5) on PL2 graph in early rounds. D2-0.0 and D2-0.5 schemes result in lower relative errors than Baseline and Bloom Filter in the first and second rounds, especially by CC and PL metrics. All schemes are comparable at $t = 3$, except on CC metric. Finally, Baseline and Bloom Filter are almost equivalent in terms of utility and they perform better D2 schemes at $t = Diam(G)$ on PL, APD and $Distance$ metrics.

VII. Conclusion

We motivate the private link exchange problem as an alternative to social graph crawling and centralized anonymization of data. The problem is distributed and provides a privacy/utility trade-off for all nodes. Our proposed problem is unique in the sense that the disseminated data over the links are the links themselves. We present two schemes for (α, β)-exchange protocol: Baseline and Bloom filter based. Experiments on synthetic graphs clarify advantages and drawbacks of both schemes. Baseline scheme keeps link lists in clear form, so its communication cost increases fast. Bloom Filter scheme incurs lower communication complexity but needs an extra recovery step in the final round. Both schemes guarantee link privacy in the range $[\frac{1}{2\beta}, \frac{1}{\beta}]$. In Baseline, the inference attack based on link counting is not much better than the random attack. For future work, we plan to investigate asynchronous models and node/links failures. We also consider community-based link exchange models in which nodes are gathered in super nodes and the link exchange takes place among super nodes only.

References

[1] Facebook Directory. https://www.facebook.com/directory.

[2] N. T. Bailey et al. *The mathematical theory of infectious diseases and its applications*. 1975.

[3] B. H. Bloom. Space/time trade-offs in hash coding with allowable errors. *Communications of the ACM*, 13(7):422–426, 1970.

[4] A. Broder and M. Mitzenmacher. Network applications of bloom filters: A survey. *Internet mathematics*, 1(4):485–509, 2004.

[5] A. Campan and T. M. Truta. A clustering approach for data and structural anonymity in social networks. In *PinKDD*, 2008.

[6] C. Dwork and A. Roth. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3-4):211–407, 2014.

[7] T. Fawcett. An introduction to roc analysis. *Pattern recognition letters*, 27(8):861–874, 2006.

[8] A. J. Ganesh, A.-M. Kermarrec, and L. Massoulié. Peer-to-peer membership management for gossip-based protocols. *Computers, IEEE Transactions on*, 52(2):139–149, 2003.

[9] G. Giakkoupis, R. Guerraoui, A. Jégou, A.-M. Kermarrec, and N. Mittal. Privacy-conscious information diffusion in social networks. In *Distributed Computing*, pages 480–496. Springer, 2015.

[10] D. A. Huffman et al. A method for the construction of minimum-redundancy codes. *Proceedings of the IRE*, 40(9):1098–1101, 1952.

[11] A. Moffat, R. M. Neal, and I. H. Witten. Arithmetic coding revisited. *ACM Transactions on Information Systems (TOIS)*, 16(3):256–294, 1998.

[12] Y. Moreno, R. Pastor-Satorras, and A. Vespignani. Epidemic outbreaks in complex heterogeneous networks. *The European Physical Journal B-Condensed Matter and Complex Systems*, 26(4):521–529, 2002.

[13] M. E. Newman. The structure and function of complex networks. *SIAM review*, 45(2):167–256, 2003.

[14] H. H. Nguyen, A. Imine, and M. Rusinowitch. Anonymizing social graphs via uncertainty semantics. In *ASIACCS*, pages 495–506. ACM, 2015.

[15] R. Pastor-Satorras and A. Vespignani. Epidemic spreading in scale-free networks. *Physical review letters*, 86(14):3200, 2001.

[16] R. Shokri, P. Pedarsani, G. Theodorakopoulos, and J.-P. Hubaux. Preserving privacy in collaborative filtering through distributed aggregation of offline profiles. In *RecSys*, pages 157–164. ACM, 2009.

[17] T. Tassa and D. J. Cohen. Anonymization of centralized and distributed social networks by sequential clustering. *TKDE*, 25(2):311–324, 2013.

[18] S. Voulgaris, D. Gavidia, and M. Van Steen. Cyclon: Inexpensive membership management for unstructured p2p overlays. *Journal of Network and Systems Management*, 13(2):197–217, 2005.